sidelobe, can also drag the DF bearing estimate, even though there is no danger of selecting this secondary cluster as the bearing estimate (if our averaging is good). Thus, as well as designing an array with a narrow beamwidth, we would like to minimise the sidelobes.

The final constraint we have is that we would like to minimise the number of array elements, as this has a cost impact. Thus, we would like a large aperture, but with few elements and low sidelobes. These are in general conflicting requirements, and so we must find some suitable compromise. In order to maximise aperture with a minimal number of elements we can choose a large aperture array with many closely-spaced elements (a filled array), and then thin it by removing some of the elements. However, this has the undesirable effect of increasing the sidelobe level compared to the filled array.

So our preferred DF array configuration is to use a thinned minimum redundancy spacing, with elements at 0, 1.0, 2.5, and 3.0λ, as illustrated at 62 in Figure 7. Figure 7 also shows the example beam patterns 70 for this configuration, and compares them to the beam patterns 72 for a conventional equally-spaced linear array of the same number of elements.

Doppler in a Non-Uniformly-Spaced Array

Another problem with the 'drawing slant lines' technique of Doppler frequency correction mentioned above (Method 3) occurs when we have non-uniform antenna element spacing. In this case, the effect of Doppler or carrier-induced frequency offset is not quite so straightforward, as it leads to a 'warping' or distortion of the DF function instead of the straightforward bearing-shift of Figure 4.

To correct for the Doppler for each time-domain DFT bin (phase sample) in this case we must use the unwinding technique (Methods 1 and 2), where the phase-shift of the unwinding between consecutive elements is proportional to the spacing between them. In principle, we can precisely correct for Doppler in this way. Alternatively, for all methods, when doing our spatial DFT we can zero pad for absent elements.

Accuracy Requirement in the Rural Scenario

Most of the above discussion has revolved around the urban scenario where, as indicated above, our error budget allows for 4.8° rms bearing error. However, in the large-cell rural environment, where the cell radius may be as high as 10km, we must achieve an angular accuracy of 0.48° rms error to meet the FCC E911 mandate. Applying equation (1) this indicates that we need to achieve a phase accuracy of around 9° between the end elements of a 3λ array. This could prove very challenging, especially when we take into account the effects of phase errors due to filters and cables, as well as the effects of noise and interference.

It is proposed that in the large-cell-radius rural scenario we increase the 'effective' array aperture to around 20λ (3 metres at 1900 MHz) as illustrated in Figure 8. We do this by deploying two 3λ DF arrays 100, accurately aligned to point in identical directions, but with a 20λ separation between their phase centres. Two single elements spaced by 20 wavelengths would produce ambiguous indications of the position of an MS 102 every 3°. However because we are using two arrays, rather than single elements, we can avoid these ambiguities. This is because two arrays of the type proposed, used together as

a pair, would ensure an accuracy of at least 1.5° between them. This would then resolve the 360° ambiguities of the high accuracy measurement achieved by comparing the mean phase of the widely spaced sub arrays. In effect, we can use each array individually to give a coarse angle estimate. We then look at the difference in mean phase between the individual arrays to give a fine angle estimate. This fine angle estimate is highly ambiguous, as there are many possible values it can take, all spaced by about 3°. However, the coarse estimates tell us which one of the ambiguous fine estimates is the correct one.

Another way of looking at this is to consider the combined beam pattern of the composite two sub-array system. In Figure 8 we show a plot 104 illustrating that the composite (boresight) beam pattern has a very narrow central main lobe. The ambiguous lobes are due to the widely spaced arrays suppressed by the inherent beam pattern of each sub-array. Assuming accurate alignment and calibration, such an arrangement will give us an extremely accurate bearing estimate, since bearing estimation accuracy is a direct function of main lobe width. We are not too concerned about the high sidelobes of such an arrangement, since in the rural scenario we expect the scattering to be benign. However, these sidelobes rule out this configuration as being unsuitable for increasing angular accuracy in the urban environment.

The invention is not limited to the embodiment of two 3λ arrays separated by 20λ. Many other combinations of dimensions are possible, as long as the widths of the individual arrays are large enough to give sufficient bearing location accuracy to remove the ambiguities due to the wider spacing between the arrays.

Limits due to array calibration

Even with plane-wave sources, high CNR and zero mobile Doppler, we can still make errored bearing estimates due to imperfections in our DF equipment if these imperfections are not perfectly calibrated out. Some possible sources of such error are:

- Differential phase errors in filters

- Phase errors at array caused by adverse weather conditions (e.g. snow build-up on radome)

- Incorrect alignment of antenna arrays

- Antenna tower twisting, due to the effect of wind, temperature cycling, etc.

In principle, none of these phenomena need cause any bearing error, as long as they can be correctly calibrated out at sufficiently frequent intervals. Naturally, it highly desirable that this can be done wholly automatically without the need for human intervention, as this would lead to much reduced network operation costs.

Embodiments of the invention provide a number of techniques for carrying out this calibration. For measuring filter/cable differential phase shifts we propose injecting a calibration signal into the receiver chain by coupling in a calibration signal close to each antenna element. The signal should preferably be injected as close as possible to each antenna element. This signal can be spread or non-spread, and can be tuned to any wanted carrier frequency in the received band. The principle benefit of using such a calibration source is

that it is unobtrusive - that is, it will not affect any other part of the CDMA BTS or network. It is also visually unobtrusive.

It would be desirable to be able to apply a similar source such that it could also calibrate out any phase errors due to weathering or ageing. To do this we would require a near-field probe (transmitter), for example extended forward from the antenna facet on a support arm, such that the probe appears within the directional pattern of the individual antenna elements. Such a scheme may adversely affect the visual impact of the antenna facet. It would also be unable to account for any errors due to antenna misalignment or tower twisting.

To account for these latter two error effects (as well as other error effects), two possible techniques can be applied according to further embodiments of the invention. The first is an Autonomous Beacon Mobile (ABM). This is in essence a standard IS95 mobile handset, mounted remotely, and provided with a long-duration power source (e.g. solar powered with battery back-up). The ABM needs to be mounted at a known (i.e. accurately surveyed) location with benign scattering (e.g. on the top of a prominent high building or mast, clear of local scattering). Alternatively, the ABM can be mounted indoors, but with an external antenna which meets the above criteria. In order to reduce the effect of scattering, the ABM could use a highly directional antenna. However, this may be undesirable in certain applications because it would restrict the ABM to being used to calibrate only a single DF BTS. More desirable would be to place the ABM somewhere in a 3-way (or more) handoff area, using an omnidirectional antenna. This way it can be used to calibrate multiple BTSs. The ABM could

also be used as an autonomous calibration tool for competitor position-location techniques such as time difference of arrival (TDOA), if such regular calibration should be needed.

To carry out a calibration, a call to the ABM is set up. The ABM incorporates circuitry to answer automatically, and to set up a dummy call. Whilst this call is in progress the DF processor at the BTS carries out a position location estimate, evaluating either or both of the bearing of the ABM and the round trip delay (RTD) from the BTS to the ABM and back. Any difference between the estimated position and true position indicates errors in the DF receiving equipment, which can then be calibrated out using a suitable calibrator. For example, measured errors in relative phases across the array can be stored in a lookup table and used later to make calibration corrections for true emergency calls. Excessive errors are an indication of equipment fault, and can be used to alert maintenance personnel.

The main benefit of the ABM is that it is simple and non-intrusive, since calls placed to the ABM provide minimal interference to the rest of the network, slightly reducing the system capacity merely for the duration of the call (i.e. only a few seconds). Of course, calls to the ABM could be restricted to off-peak hours. The main disadvantage of using ABMs is that the operator needs to obtain sites at which to locate them. If we are able to place every ABM in (or close to) a 3-way handoff area, then we will need an average of one ABM site for every BTS cellsite (for a trisector system).

A solution to this site-acquisition problem is to place beacons actually at the cellsites themselves. Of course,

such a beacon cannot be used to calibrate its own cell, but it *can* be used to calibrate the surrounding cell sites, if it has sufficient transmit power. An added advantage is that the scattering environment is almost guaranteed to be benign, since cell masts are generally located clear of local scatterers, and the range will be high since there should be good line of sight (LOS) coverage to surrounding sites. However, the ABM concept cannot necessarily simply be extended to positioning an ABM at each cellsite. For example, if we used a ABM located at a cellsite in a CDMA system, then in order to avoid interfering with its own cellsite the ABM would be power-controlled down to extremely low power levels. It would then be undetectable at the surrounding cellsites. Therefore, in order to be able to use a cellsite beacon (CSB) with sufficient power we need to give it a carrier frequency within the CDMA band (for licensing reasons, and so that it can be detected by the DF systems it is being used to calibrate), but away from all of the active channels being used by that cellsite. In a preferred embodiment the CSB therefore transmits a narrowband carrier within one of the operator's guard-bands (i.e. at the edge of the licensed allocated Base Station Receive band). This is proposed in terms of a CDMA system but may easily be extended to other systems.

How powerful can we make this narrowband CSB carrier? The main restriction on CSB transmit power in a CDMA system is identified to be due to the specification for adjacent-channel single-tone desensitisation of a BTS receiver (as discussed in more detail below). To avoid desensitising the BTS receiver it is proposed that we restrict the tone transmit power to -12dBm. Based on a 10Hz receiver bandwidth at the DF receiver to be calibrated (i.e. coherent averaging of the received